SANDIA REPORT

SAND98–2718
Unlimited Release
Printed January 1999

Lessons Learned in Interconnecting Sandia's Network Backbone Routers to a Core ATM Switch

Thomas V. Archuleta

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy under Contract DE-AC04-94AL85000.

Approved for public release; further dissemination unlimited.

Sandia National Laboratories

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Prices available from (703) 605-6000
Web site: http://www.ntis.gov/ordering.htm

Available to the public from
National Technical Information Service
U.S. Department of Commerce
5285 Port Royal Rd
Springfield, VA 22161

NTIS price codes
Printed copy: A03
Microfiche copy: A01

SAND98-2718
Unlimited Release
Printed January 1999

Lessons Learned in Interconnecting Sandia's Network Backbone Routers to a Core ATM Switch

Thomas V. Archuleta
Telecommunications Operations I
Sandia National Laboratories
P.O. Box 5800
Albuquerque, NM 87185-0812

Abstract

We learned some valuable lessons during the migration of the corporate unclassified production network's Cisco Systems 7000 series routers to the Lucent Technologies GlobeView 2000 ATM switch. This paper documents those lessons learned and discusses in some detail each phase of the migration process. The unclassified production network at Sandia National Laboratories consists of various telecommunications devices. Some of the devices include Cisco routers, Cisco Asynchronous Transfer Mode (ATM) and Ethernet switches, and Bay Networks hubs. The Cisco 7000 series routers and some LightStream 1010 (LS1010) ATM switches are interconnected with the GlobeView ATM switch at the core of the network. Prior to introducing ATM into the production network, the Cisco routers were interconnected on a Fiber Distributed Data Interface (FDDI) ring at the core with High Speed Serial Interface (HSSI) point-to-point connections to remote routers in various geographic locations throughout Sandia's campus.

Table of Contents

List of Figures

Intentionally Left Blank

Introduction

When we migrated the corporate production network's Cisco 7000 series routers to the GlobeView ATM switch we developed some key and learned some valuable lessons. This paper begins by describing Sandia's production network and the GlobeView ATM switch. It then discusses the migration phases in detail with supporting diagrams to show the documentation methods used in planning the migration process. Finally, this paper describes the techniques used to support the migration process and presents the lessons we learned as we moved the production routers from FDDI connectivity to a core ATM switch.

The production network at Sandia National Laboratories consists of various telecommunications devices. Some of the devices include Cisco routers, Cisco ATM and Ethernet switches, and Bay Networks hubs. Several Cisco 7000 series routers are interconnected with the GlobeView 2000 ATM switch at the core of the network. We moved interconnection of the routers from the FDDI configuration to the GlobeView ATM switch in a series of phases. We began initially by building a testbed consisting of three routers interconnected with an ATM Permanent Virtual Circuit (PVC) mesh through the GlobeView and tested interoperability where possible on all network protocols. These protocols included Transport Control Protocol/Internet Protocol (TCP/IP), AppleTalk (AT), Internet Packet eXchange (IPX), Digital Equipment Corporation NETwork (DECNET), Xerox Network Service (XNS) and ConnectionLess Network Service (CLNS). As would be expected, trying to build client/server functionality for the entire set of network protocols is virtually impossible. However, we reasonably tested interoperability of TCP/IP, AT, IPX and DECNET as these are the primary protocols supported on the production network.

As we planned the phases of migration, we identified several techniques to support our efforts. One of these techniques included a configuration scheme where a device could be recognized by its address on the network regardless of the protocol. This configuration scheme provided documentation of the network within the network itself and supports ease of troubleshooting. As we implemented the phases of the migration process, we ran into some interesting problems and learned some valuable lessons. We applied those lessons to the migration process as we progressed, resulting in alteration of our implementation plan. The new plan entailed pre-configuring and pre-installing Cisco 7000 routers and moving customer circuits to the new configuration in a series of smaller phases. This minimized the network downtime seen by our customers.

As you read through this document, you will venture through brief descriptions of Sandia's production network, the Network/Configuration Management applications used at Sandia, and the GlobeView ATM switch. You will see a detailed discussion of the migration phases with a presentation of the key techniques we used and the valuable lessons we learned. Finally, there is summary of those lessons learned and an appendix providing illustrations used in supporting the migration process.

Description of the Open and Restricted Networks in 1995

Sandia's networks are comprised of Metropolitan Area Networks (MAN, i.e., the two main campuses in Albuquerque NM and Livermore CA), Wide Area Networks (WAN, i.e., inter-campus and external collaborations), and Local Area Networks (LAN, i.e., intra-building communications). At Sandia, we have three network environments. One is Open/Internet, the second is Restricted (i.e., Intranet behind firewalls), and the third is Classified. The Classified network configuration was not subject to migration to the GlobeView. Therefore, this section describes the configuration of the Open and Restricted Networks as they existed in 1995 prior to the migration process.

We have Cisco 7000 series routers located in major buildings at the Albuquerque campus. These routers provide network backbone connectivity for the intra-building devices including LAN switches, FDDI concentrators and Ethernet hubs. In the main facility, we have core 7000 series routers, which were interconnected on three FDDI rings. In Figure 1-1, the Restricted Network consisted of routers connected to a FDDI ring with routers in major buildings interconnected to the core using High Speed Serial Interface (HSSI) point-to-point links. The dotted lines in Figure 1-1 represent links using Synchronous Optical NETwork (SONET) MUXes for connectivity. The Restricted Network resides behind a firewall system, which included a router with multiple FDDI interfaces. The External and Restricted Network configurations are similar in using FDDI and HSSI for interconnection between routers. The OutNet Network consists of routers provided by our Internet Service Providers. At the remote WAN sites, we use Cisco 2500 series routers (not shown in Figure 1.1) with connectivity to a core 7000 router located in the main facility. We also have a "private" line to Livermore, CA on each network.

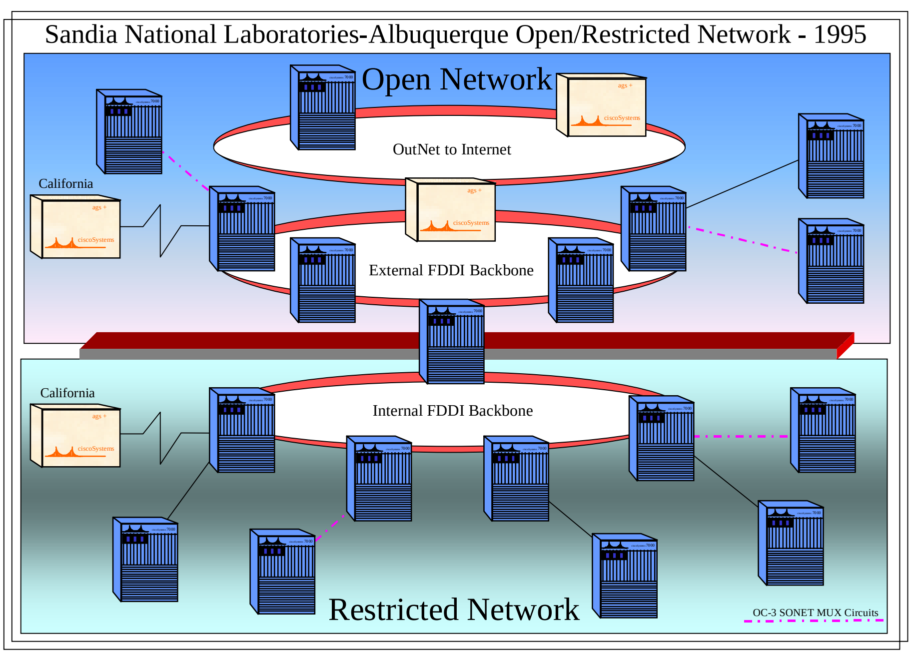

Figure 1-1. Diagram of Network Core – Last Quarter 1995.

Other technologies utilized in the network core include SONET/DS-3 MUXes, General Datacom Corporation (GDC) MUXes and Channel Service Unit/Data Service Units (CSU/DSU). Prior to the migration process, we supported the major protocols TCP/IP, AppleTalk, IPX, DECNET, XNS, and CLNS. As we began the migration process, we eliminated CLNS, as it was no longer needed. Refer to Appendix A-1.1 for an illustration of the Internal Restricted Network (IRN) Backbone original configuration.

As you can see, Sandia maintains a complex network configuration utilizing a variety of telecommunications devices. The IRN and External Open Network (EON) contained the same basic configuration with a router performing packet filtering as part of a firewall system protecting the IRN. The connectivity from remote buildings within the main campus to the core network backbone utilized the HSSI links.

Network/Configuration Management

Sandia's Telecommunications Operations Department I network operations staff is committed to using configuration and network management utilities to mange the corporate production network. We perform the Cisco device configuration management from Sun Microsystems workstations running management software from Cisco Systems. Our network management utilities poll the network devices 24 hours per day, 7 days per week at 3-minute intervals. The following section briefly describes the applications we use for configuration and network management on the production network.

A Sun Ultra-60 SPARCstation running Solaris 2.6 Operating System software supports our configuration management efforts on the IRN. We currently use SunNet Manager Network Management System (NMS) on the Ultra-60 as a base for the CiscoWorks 4.0 (CW) suite of applications. CiscoWorks 4.0 and CiscoWorks 2000 compose the configuration management function for the Cisco devices (Routers, ATM and Ethernet switches.) Cisco currently supports managing Virtual LANs in an enterprise with CiscoWorks for Switched Internetworks (CWSI) Campus 2.2.

The CWSI Campus 2.2 application includes tools such as CiscoView, VLAN Director, ATM Director, and Traffic Director. Refer to Appendix A-2.1 through A-2.4 for illustrations of the CiscoView, VLAN Director, ATM Director, and Traffic Director applications.

Figure 2-1 illustrates the network mapping supported by SunNet Manager. The map in Figure 2-1 mainly represents the Cisco Routers, ATM switches, and Ethernet switches within building 880. Note the “Tools” window within the SunNet Manager shows the menu containing the CiscoWorks applications.

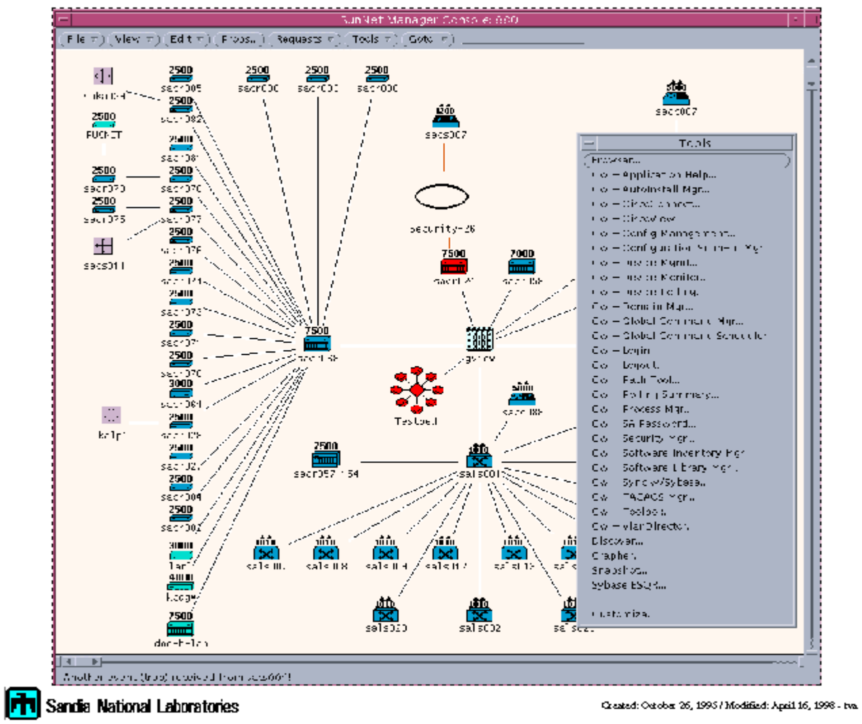

Figure 2-1. SunNet Manager and CiscoWorks 4.0 Applications.

The CiscoWorks 2000 management software provides enhanced management capabilities, for networks using ATM and VLAN technology. Refer to Figure 2-2 on the next page.

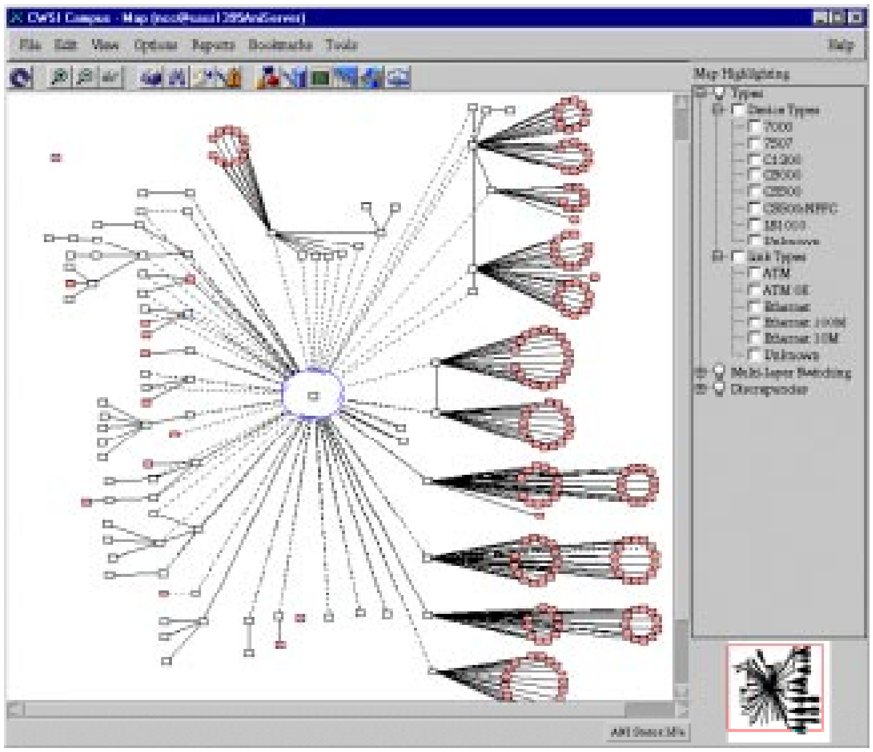

Figure 2-2. CiscoWorks 2000, CWSI Campus 2.2 Map of the IRN

Figure 2-2 shows the CWSI Campus 2.2 map of the IRN Virtual Network. If you compare the SunNet Manager map from Figure 2-1 with the CWSI Campus 2.2 map, you will see a significant difference in the mapping of objects (network devices). This is because of the manual vs. automatically discovered placement of the objects on the map by the two different applications. Cisco also provides a web-based management system as part of the CW 2000 package called Resource Manager Essentials (RME). Refer to Appendix A-2.5 for an illustration of the RME display.

We use the Seagate Enterprise Management Software DiMONS to poll network devices on a regular basis (24x7). We use DiMONS primarily for OPERATIONAL status polling. We normally do not poll network devices from the configuration management station unless necessary in troubleshooting. We also use Bay Networks Optivity Suite for hub management. We are in the process of migrating these network management functions to be supported by Hewlett-Packard Network Node Manager OpenView with NerveCenter (formally DiMONS). The HP OpenView NMS supports integration of the Cisco and Bay applications to provide network/configuration management of ATM, VLAN, and legacy Ethernet network technologies.

Testing ATM Interconnectivity

As we began planning the migration process, we utilized the test bed to verify the interoperability of Cisco 7000 routers interconnected through the GlobeView ATM switch. Following is a brief description of the testing we performed and the future of ATM in the production network backbone.

We connected three Cisco 7000 routers in a fully connected PVC mesh on the GlobeView ATM switch. Two routers were configured on the same TCP/IP subnetwork to test protocol routing functionality, as it would be configured on the production ATM network backbone. The third router provided connectivity between the test bed and the production network. We used TCP/IP services such as telnet, ftp, etc., to verify the interoperability of the Cisco routers in the PVC mesh on the GlobeView. Although it is rather difficult to simulate the production environment in a small test bed, the tests proved successful. We also tested all other protocols supported on the production network to the best of our ability.

While migrating to a core ATM network backbone, Sandia's Advanced Networking Integration Organization analyzed the functionality of LAN Emulation (LANE) over ATM on the test bed to support legacy Ethernet network technology in the production network. The test bed was expanded to incorporate Cisco LS1010 ATM switches and Catalyst 5000/5500 Ethernet switches to support LANE. This network technology has been implemented on the production network backbone. We also started a "user-trials" configuration in the test bed, where some of the customers requiring higher bandwidth were using the test bed. This configuration currently exists on the production network backbone along with LAN Emulation.

The GlobeView ATM Switch

This section provides an overview of the Lucent Technologies GlobeView 2000 ATM Switch. Sandia's version of the GlobeView is Release 1.5. The GlobeView has a 20 gigabits per second switching fabric, with a fully redundant architecture. It is composed of five cabinets, two contain the control processors, switching fabric, and Interface and Services Stage (ISS) Units, three other cabinets contain even more ISS shelves.

The GlobeView supports a maximum of 128 OC-3 ports. Sandia's GlobeView is configured with mostly OC-3 interfaces (approximately 50 ports); however, we have some DS-3 interfaces as well. Figure 3-1 shows a rear view of the GlobeView.

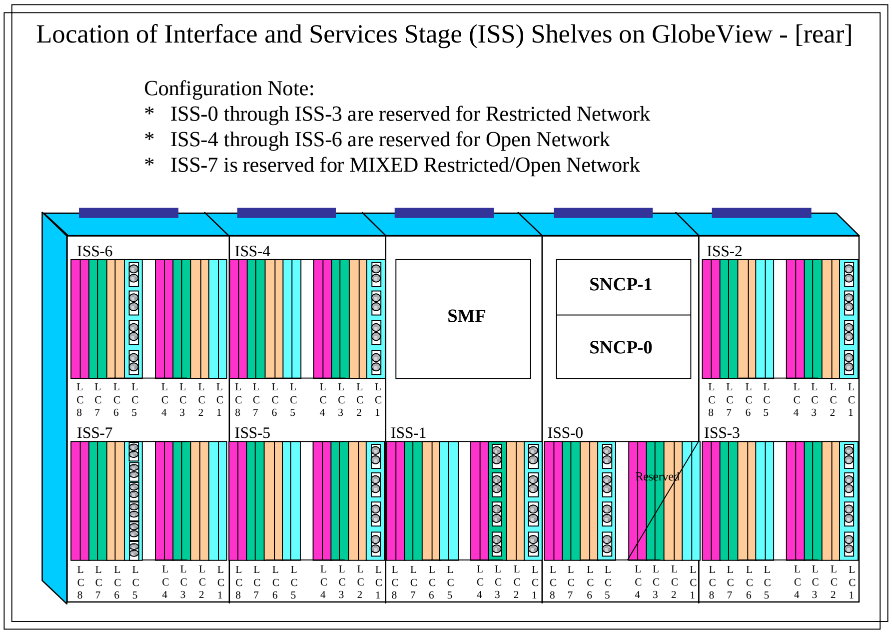

Figure 3-1. Rear View of the Lucent Technologies GlobeView 2000 ATM Switch.

Note the location of the SNCPs (Service Node Control Processor) and SMF (Shared Memory Fabric). Each ISS cabinet consists of two ISS shelves with two groups of control modules for each side. We have chosen ISS-0 through ISS-3 for the Restricted Network (ISS-0 is reserved for the test bed) and ISS-4 through ISS-6 for the External Network (ISS-6 reserved for the test bed). The ISS-7 shelf supports a mixed environment, Restricted and Open. The intent of this shelf is to support remote connections, i.e. the “private” line between SNL/Albuquerque, NM and SNL/Livermore, CA, referred to the Intersite Link throughout this paper. This intersite link utilizes separate ATM PVCs to provide network connectivity between Sandia’s two main campuses for both Restricted and Open networks.

The Migration Process

This section provides detailed discussion of the migration process used to move the routers from the FDDI interconnectivity to the GlobeView ATM switch. The discussion for each phase consists of a diagram illustrating the planned migration for the particular phase, the key techniques used, and the lessons learned during the appropriate phase.

Our migration plan initially began as four phases. Phases 1-3 completing the IRN and Phase 4 completing the Intersite Link. We had not yet considered the EON as of this time. In Figure 4-1, you will see the number of routers planned for each phase.

The Migration/Cutover Plan

Initially - Four Phases:
Phase 1: Internal Restricted Network (IRN) - 5 routers
Phase 2: IRN - 5 routers
Phase 3: IRN - 8 routers
Phase 4: Intersite Link

Actual - Five Phases:
Phase 1: IRN - 5 routers
Phase 2: IRN - 3 routers
  Problems noted in Lessons Learned
Phase 3: IRN - 8 routers
  Series of smaller "cutover" phases
Phase 4: Intersite Link
Phase 5: External Open Network (EON) - 16 routers
  Series of smaller phases

Figure 4-1. Number of Routers in Migration Planning, Initial vs. Actual Phases.

We set our ambitions high as we initially planned the migration process. The migration of the routers to the GlobeView took five phases with Phase-5 completing the EON. We encountered some interesting problems and learned some valuable lessons, which led us to alter our initial plan. We determined migrating the routers with smaller cutover phases was easier for us to plan, configure, verify, move, and re-verify network connectivity.

The Migration Process: Phase-1 Internal Restricted Network

This section discusses the first phase of the migration process, presents the key techniques used as we began, and lists the lesson learned during Phase-1. Refer to Figure 4-2 for an illustration of our Phase-1 migration diagram.

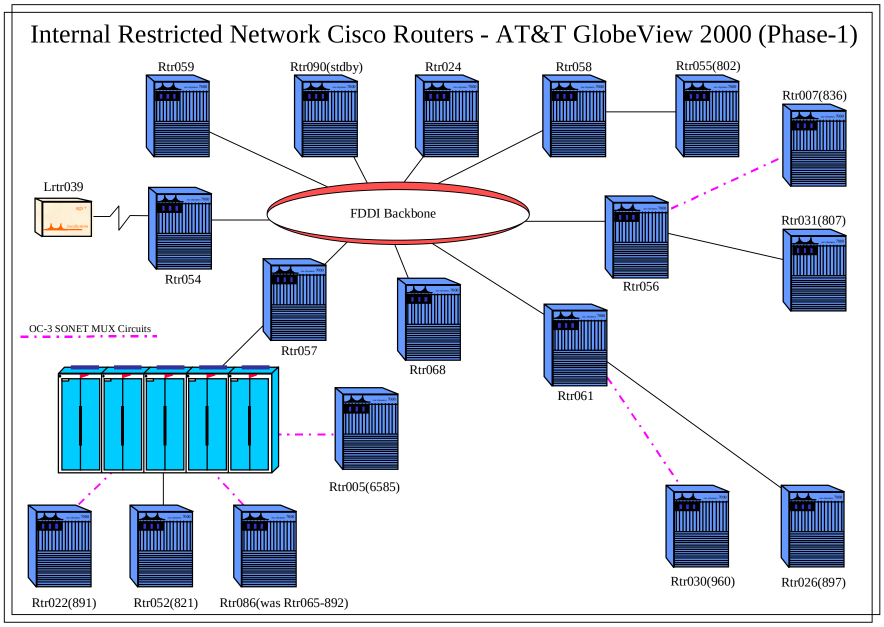

Figure 4-2. Diagram of Phase-1 of the Migration Process.

We decided to have Rtr057 as the "connection between two technologies" providing routing from the central FDDI backbone to four of the remote Cisco 7000 routers connected to the ATM switch. The Phase-1 routers were Rtr005, Rtr022, Rtr052, Rtr057, and Rtr086. These were originally DS-3 connections utilizing High Speed Serial Interfaces (HSSI) on the routers and CSU/DSUs.

The Migration Process: Phase-1 Key Techniques

We utilized various techniques for documenting the implementation of the PVCs on the router's ATM interface. The next few pages provide some definitions of ATM terms and present those techniques used in Phase-1.

Recall the following definitions from the ATM forum.[1] A VC (Virtual Circuit) is a communications channel that provides for the sequential unidirectional transport of ATM cells. A VP (Virtual Path) is a unidirectional logical association or bundle of VCs. A VCI (Virtual Channel Identifier) is a unique numerical tag as defined by a 16 bit field in the

[1] The ATM Forum Web-site, Glossary: http://atmforum.com/atmforum/library/glossary/gloss-v.html

ATM cell header that identifies a virtual channel, over which the cell is to travel. A VPI (Virtual Path Identifier) is an eight-bit field in the ATM cell header, which indicates the virtual path over which the cell should be routed.

In researching the configuration requirements for the router AIP (ATM Interface Processor) from Cisco Systems[2], note the definition of VCD (Virtual Circuit Descriptor) as a unique number per AIP that identifies to the AIP which VPI/VCI to use for a particular packet. The AIP requires this feature to manage packet transmission.

We map the VCD to the remote router's host address for IP protocol, since the router uses the VCD to keep track of the VP/VC pairs to send traffic. Figure 4-3 illustrates our initial mapping of router host addresses.

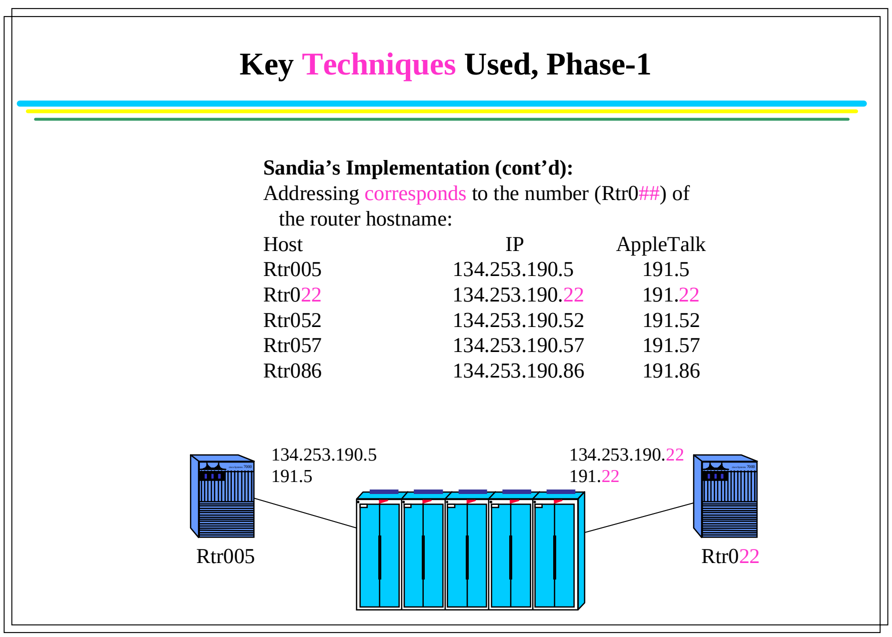

Figure 4-3. Phase-1 Router Addressing.

The addressing of the routers corresponds to the Rtr0## numbers in the hostnames. The router Rtr022 has IP address 134.253.190.22 and AppleTalk address 191.22. This method will document the network from within the configuration and provide address mapping when someone is deep into troubleshooting by using device "show" commands.

The GlobeView uses the VPI, VCI, and physical port to provision the circuit. Figure 4-4 illustrates provisioned circuits for the Phase-1 router PVC mesh.

[2] Cisco Systems Web-site, ATM Reference Guide: http://www.cisco.com/univercd/cc/td/doc/product/software/ios103/rpcs/78788.htm

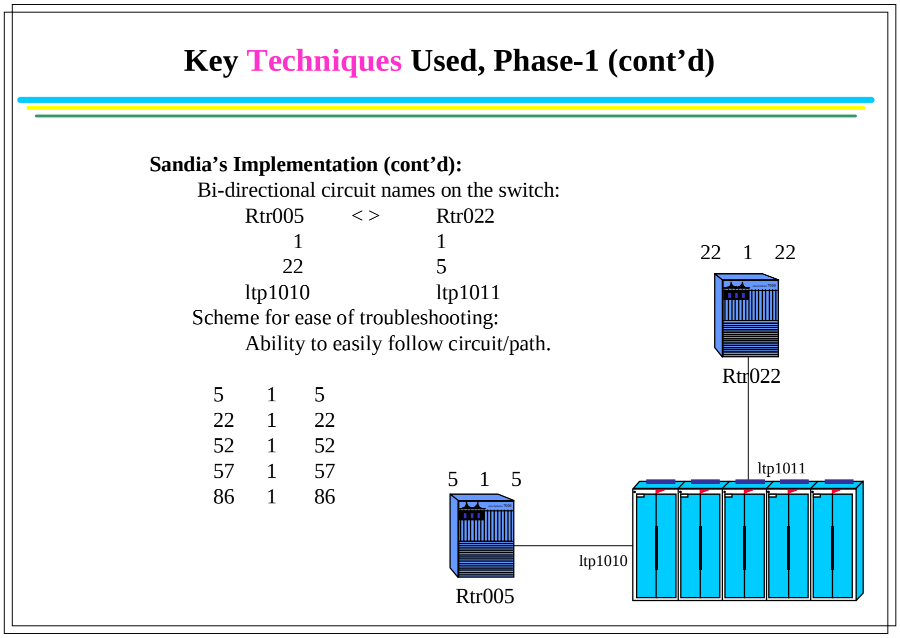

Figure 4-4. GlobeView Provisioning for Phase-1 Router PVC Mesh.

The GlobeView supports provisioning of bi-directional circuits, therefore our naming convention used the < > symbols to indicate the bi-directional circuit. The circuit configuration includes VPI=1. The VCI points to the other router's port on the GlobeView. This port is called the Logical Termination Point (LTP) and ltp1010 represents the actual port where Rtr005 is connected. We chose to configure the VCD and VCI to be the same because of the method by which the routers and GlobeView "build" the circuit. This indeed provided ease in understanding the configuration of the PVC mesh because all addressing and VPI/VCI numbers are based on the fourth octet of the router's IP address.

The remaining key techniques used in Phase-1 include: building the "cutover" router's configuration in advance, utilizing ISDN connectivity via kermit on a Sun workstation to each router's console port, and utilizing the benefits of X-Window CUT/PASTE to apply the new configuration to each router. In an effort to fully understand the physical connectivity of each router, a detailed map showing the complete circuit was created before each migration phase to support backing out of the migration if needed. A detailed map was also generated after each migration phase to provide complete documentation in the event of physical troubleshooting. Refer to Appendix 3-1 and 3-2 for mapping of the first four routers in Phase-1.

The Migration Process: Phase-1 Lessons Learned

We considered the initial Lessons Learned to be minor. Basically, three lessons were learned during Phase-1 of the migration.

The first lesson learned was the result of inaccurate IP routing protocol Open Shortest Path First (OSPF) configuration. The OSPF routing protocol is a link-state protocol.[3] If you consider a link to be an interface of a router, the state of the link is the relationship of the interface to neighboring routers. OSPF uses “flooding” to exchange link state information between neighboring routers. In order to control the flooding of route updates, OSPF uses Areas to put a boundary on the flooding of information. The FDDI connected routers all resided in the “backbone area”, i.e. Area 0. The Rtr022 router was the first router we moved directly to the ATM switch. It was located in Area 57 behind Rtr057; i.e. Rtr057 was a boundary router to Rtr022. In setting up OSPF routing on Rtr057 (recall Rtr057 is link between the ATM Switch and FDDI ring), we had extended the AREA 0 to the switch for Rtr022. As we applied the configuration to Rtr022 for OSPF, we had AREA 0 and AREA 57 crossed because of the way OSPF gets configured on the router - on a line-by-line basis. We had to re-do OSPF routing on Rtr057 in order for all IP routing problems to clear.

The second lesson learned was setting up the route caching for the ATM Interface Processor (AIP). The ip route-cache cbus (autonomous fast-switching) command had been placed on the ATM sub-interface rather than the main ATM interface. We mainly lost time tracking down the correct method of configuring the AIP for fast-switching packets to the interface.

The third lesson learned resulted in undesired network change to customer LANs. We had asked the customers affected by the routers which were migrated to assist us in verifying a successful cutover. Some of these customers decided to take advantage of the production network downtime and made changes in their LANs. During the verification process, we found problems with network connectivity, which eventually resulted from their changes. This only made extra troubleshooting work for us. *Do not let any other network changes occur or control all network changes!*

Phase-1 of the migration was completed February 10, 1996, and was somewhat problem-free. The techniques discussed in this section were significant contribution to the success of this first phase. We will see in the next section; however, that our configuration scheme was not quite adequate enough for a completely successful migration process.

[3] Cisco Connection Online Web-site, OSPF Design: http://www.cisco.com/warp/public/104/1.html

The Migration Process: Phase-2 Internal Restricted Network

In the second phase of the migration process, we planned to migrate the remaining routers connected to the production network backbone via HSSI point-to-point connections. As we progressed through this phase, we ran into some problems. This section will discuss in detail those problems and highlight the corrective action necessary to continue the migration process. Any key techniques used will be presented followed by another list of those lessons learned during Phase-2.

The Phase-2 routers were Rtr007, Rtr026, Rtr030, Rtr031, and Rtr055. (Note all routers connected to the GlobeView have building numbers in parentheses and the Central FDDI backbone only remains.) Figure 4-5 illustrates the proposed migration plan.

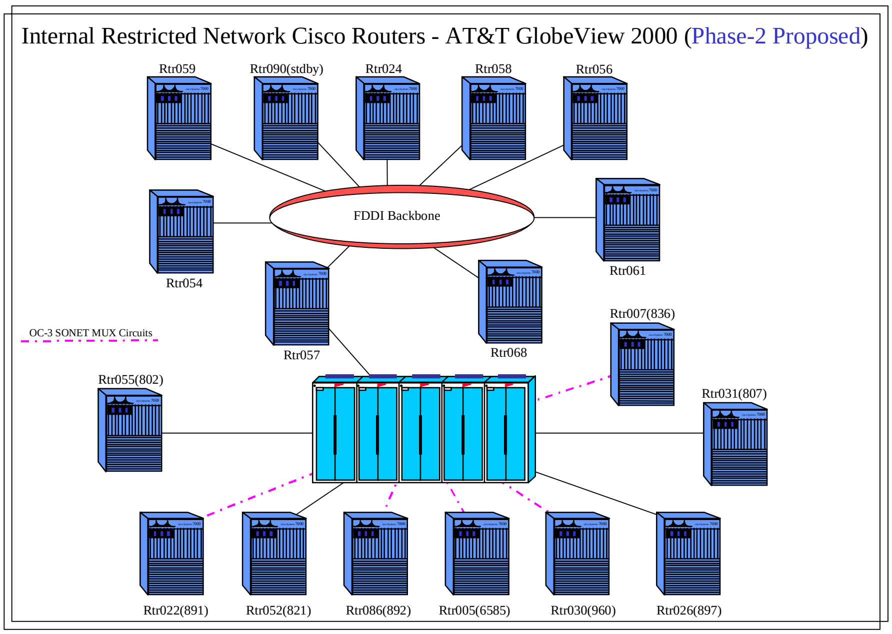

Figure 4-5. Diagram of Phase-2 of the Migration Process.

The routers connected to the FDDI ring would only remain to be moved during the next phase of the migration process.

The Migration Process: Phase-2 Key Techniques

Only one additional technique was utilized during Phase-2. In thinking about how the IPX address gets assigned to the router, we realized we might have problems in the future. See Figure 4-6.

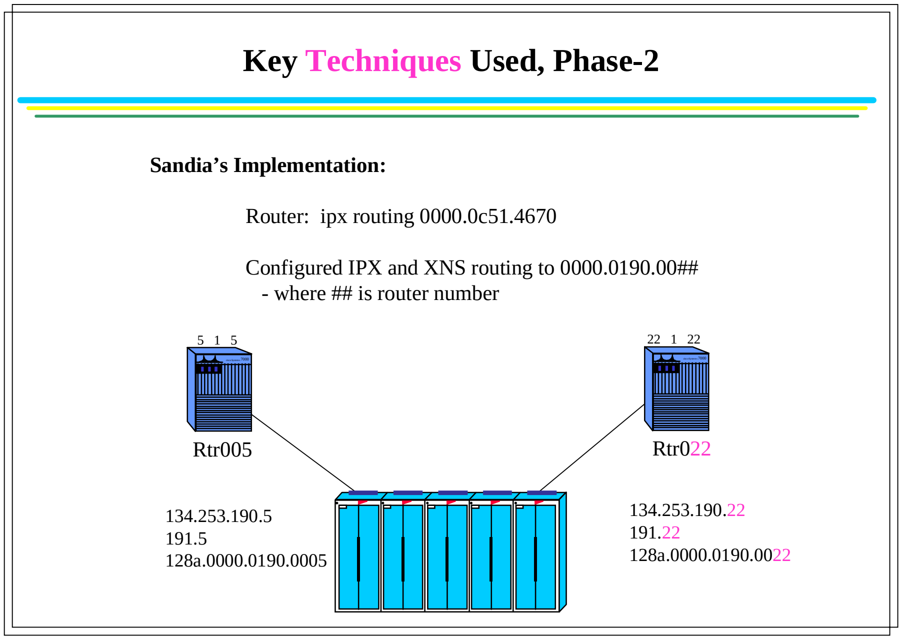

Figure 4-6. IPX Routing Configuration.

Basically, the hardware address of the first addressable interface (i.e. Ethernet, FDDI) on the router is used in creating IPX network routing. This address, along with the ATM interface IPX network address, must be assigned to a VCD in order for IPX between two routers to work. In Figure 4-6, the IPX address for Rtr022 is 128a.0000.0190.0022, where 128a is the network address and 0000.0190.0022 replaces the hardware address. In the event of a failure on the interface from which the original hardware address was obtained, we would have to re-do each router's IPX VCD map-list entry to include the new hardware address. We tried to pre-assign this address on the router with a static address to avoid having to re-do each router's map-list. The 0000.0190.00## fits the format of a hardware address,190 is the IP subnet used on the GlobeView and 0022 represents Rtr022. Note the consistency in our addressing. This method has also proven helpful in troubleshooting problems.

The Migration Process: Phase-2 Problems Encountered

During Phase-2, we encountered some problems causing delay in the migration and change in our method of migrating routers to the GlobeView. We had turned off the router's console logging because excessive IP OSPF and AppleTalk messages were cluttering the screen during our configuration changes. As we added the modified configuration to the routers, PVCs were failing to be created. We could not determine what was going wrong. If we would completely re-configure the ATM Interface on the router, it appeared to clear the problem. After all routers had been modified and appeared to be properly interconnected, we used CiscoWorks Configuration Management tools to "update" the configuration database. In this process, we noticed routers began losing network connectivity.   Upon going back to the routers via ISDN into the console and turning the console logging back on, we decided to re-configure Rtr055 ATM interface configuration. This included removing the PVCs, map-groups, addressing, and ATM rate-queue. We even shutdown the interface and started over. This is when we noticed certain PVCs would not create and generated the following error message:

```
%AIP-3-AIPREJCMD: Interface ATM4/0, AIP driver rejected Setup VC
command (error code 0x8000)
%ATM-3-FAILCREATEVC: ATM failed to create VC(VCD=5, VPI=1, VCI=5)
on Interfaces, (Cause of the failure: ATM4/0)
%ATM-3-FAILCREATEVC: ATM failed to create VC(VCD=7, VPI=1, VCI=7)
on Interfaces, (Cause of the failure: ATM4/0)
```

After much checking and consultation with Cisco, we determined that we were using the wrong VCI numbers. The VCI ranges 0-31 are reserved for ATM switch maintenance and administration functions. We had tested our VCD / VCI scheme in the test bed using VPI 1 and found no problems. Our perception of only VPI 0 reserving VCI 0-31 was wrong. To correct this problem, we simply added 100 to the VCD / VCI numbers. This was a time consuming process of correction, because each router and switch circuit on the mesh required reconfiguration.

Once the reconfiguration was complete, we still encountered a problem where PVCs would "disappear". We eventually determined, with further consultation from Cisco, that there is a bug in the router AIP microcode. We performed all tests with Cisco Internet Operating System (IOS) Enterprise 10.2(8) which contains bundled AIP microcode AIP MC 10.8. Cisco identified the bug and asked us to use AIP MC 10.11. By this time we had already backed out two of the routers to their original configuration, including one from Phase-1.

We put the AIP MC 10.11 image into the router's flash-memory. In trying to upgrade the AIP microcode from flash-memory, we encountered more problems. We found that the AIP resets because of self-tests the router performs on all interfaces during a microcode reload. Our configuration management tool contained the microcode reload statements at the beginning of the router's configuration file. Since the AIP would reset as a result of the microcode reload, the configuration management tool was unable to successfully function. We also found that random Ethernet interfaces would lose line-protocol during the microcode reload. An actual manual reset of the Ethernet interface was required, via shutdown/no shutdown or clear commands on the router.

As you can imagine, this is not an acceptable method of operation for the production network backbone. We decided to try an IOS version which would keep us at or above the recommended AIP microcode level. Note the Cisco IOS provides bundled microcode within the IOS image. We decided to try IOS Version 10.3(9) because of all problems we encountered in the 10.2(X) family. All tests in the testbed were successful. However, upon trying to implement IOS 10.3(9) on the production environment, AppleTalk would not initialize on the AIP interface thus generating the above error message.

Cisco recommends using the highest revision level in any IOS family. This is the reason for our decision to try IOS 10.3(9). Because of the AppleTalk problem, we decided to try IOS 10.2(11), which was the highest during February 1996. The IOS 10.2(11) software contains the recommended AIP MC 10.11 microcode for the AIP. All tests in the testbed were successful. We began implementation of IOS 10.2(11) on the production network, and no problems were noted. We are currently running IOS 10.2(11) on the production network. We were able to continue the migration and completion of Phase-2, including the routers which we had backed-out (Rtr026 and Rtr086), occurred on April 3, 1996.

The Migration Process: Phase-2 Lessons Learned

The lessons learned during Phase-2 were indeed major when we upgraded the GlobeView management software just prior to starting Phase-2, we added another variable to troubleshooting all router connectivity problems because of PVC create failures and PVCs dropping. Also, our locally written program for provisioning the PVC mesh on the GlobeView resulted in us missing warning messages when we provisioned circuits using the 0-31 reserved VCI numbers. This locally written program is certainly useful, though, because of the quickness in tearing down and re-provisioning the PVC mesh circuits. *Do not disable the router console logging.* Even though the AppleTalk and OSPF messages clutter the console display, it is important to note any configuration errors on the router. We found the AIP microcode bug in AIP MC 10.8, and we also found that Cisco IOS 10.3(9) does not support AppleTalk when configured on the AIP.

Phase-2 of the migration was completed on April 3, 1996, but not without many hours of intense work. We had turned off the router’s console logging and missed crucial error messages during the configuration of the routers for interconnection to the GlobeView. We had utilized reserved ATM VCI numbers in configuring the PVC mesh. We also encountered buggy Cisco IOS software. The next phase, Phase-3, will prove successful as we apply our lessons from Phase-2 to the planning of Phase-3 and alter our plan.

The Migration Process: Phase-3 Internal Restricted Network

This section presents Phase-3 of the migration process and briefly describes the changes to our plan. We had anticipated moving the remaining Central FDDI backbone routers to the GlobeView and remove the FDDI ring. Figure 4-7 illustrates our proposed migration plan.

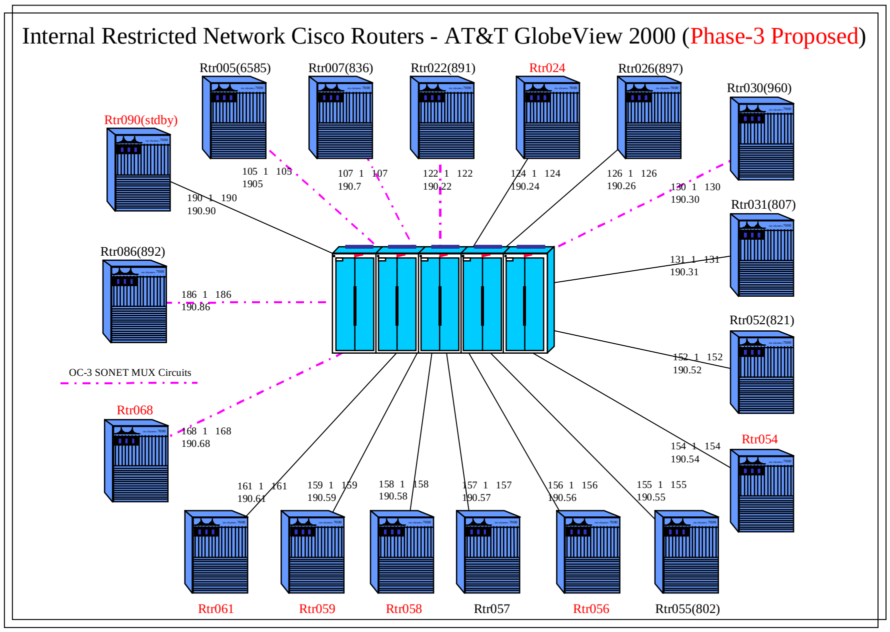

Figure 4-7. Diagram of *Proposed* Phase-3 Configuration of the Migration Process.

Recall the lessons we learned in the previous phases. A brief summary of the lessons learned so far includes: OSPF Area mismatch, incorrect AIP route-cache configuration, customers taking advantage of the network downtime, upgrading the GlobeView management software, using reserved VCI ranges 0-31, disabling router console logging, buggy AIP MC 10.8 microcode, and buggy IOS 10.3(9) AppleTalk configuration.

Because of the lessons we learned in the previous two phases, we changed our plan. Figure 4-8 illustrates our actual approach to Phase-3.

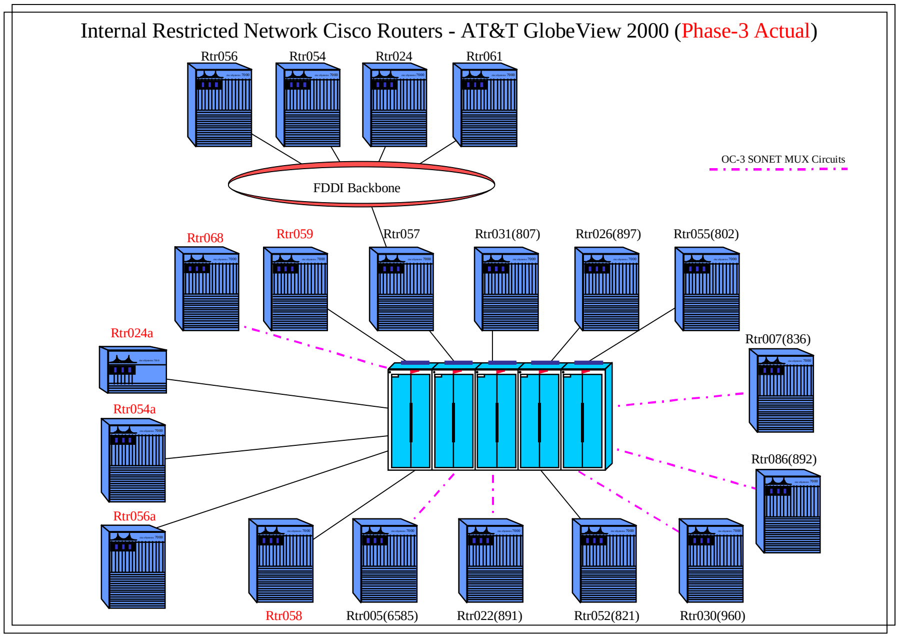

Figure 4-8. Diagram of *Actual* Phase-3 Configuration of the Migration Process.

We added (staged) three new routers to the production PVC mesh. This enabled us to move customer circuits from their FDDI router to the router connected to the GlobeView. The new routers were Rtr024a, Rtr054a, and Rtr056a in the above illustration. We still followed all previous migration phase configuration techniques and simply removed the Rtr024, Rtr054, and Rtr056 routers at a later time after moving customer circuits to the new routers. This was done in a series of smaller phases, which made the migration process much easier.

The Migration Process: Phase-3 Key Techniques

In keeping the migration process moving, the techniques identified for Phase-3 were, as mentioned earlier, to change our plan by migrating in a series of smaller phases. We also staged the new routers to replace those on the FDDI ring.

There were no lessons learned during Phase-3 of the migration.

Phase-3 of the migration was completed on May 16, 1996 thus completing the IRN. This third phase of the migration was successful and provided necessary proof of the process improvements before we continued with the Intersite Link and the EON.

The Migration Process: Phase-4 The Intersite Link

This section briefly discusses the fourth phase of the migration process. This section presents a brief description of the Intersite Link, both before and after the changes, and lists the lessons learned during Phase-4. The original configuration of Sandia's Intersite Link between Albuquerque and Livermore utilized the AT&T BNS 2000 ATM switches to complete the circuits via DS-3.

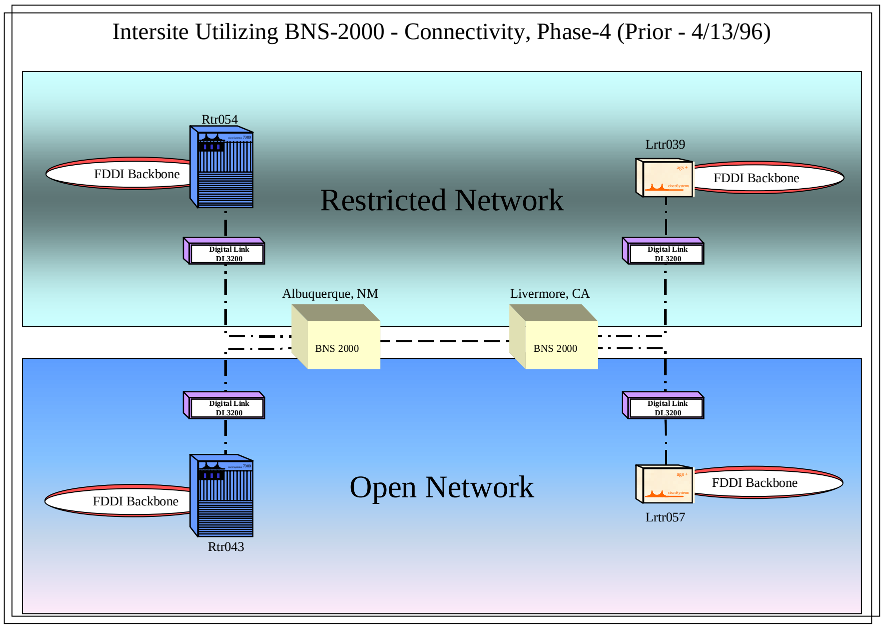

Figure 4-9. Diagram of Phase-4 of the Migration Process.

As you can see from Figure 4-9 above, this circuit utilized Digital Link DL3200s. We ran Switched Multimegabit Data Service (SMDS) over this link. We have a router in California for both the California Open Network (CON, analogous to Albuquerque EON) and Restricted Access Network (RAN, analogous to Albuquerque IRN). The interface on the router is High Speed Serial Interface (HSSI). A router on each of the Albuquerque networks (EON and IRN) supplied the routing between both campuses.

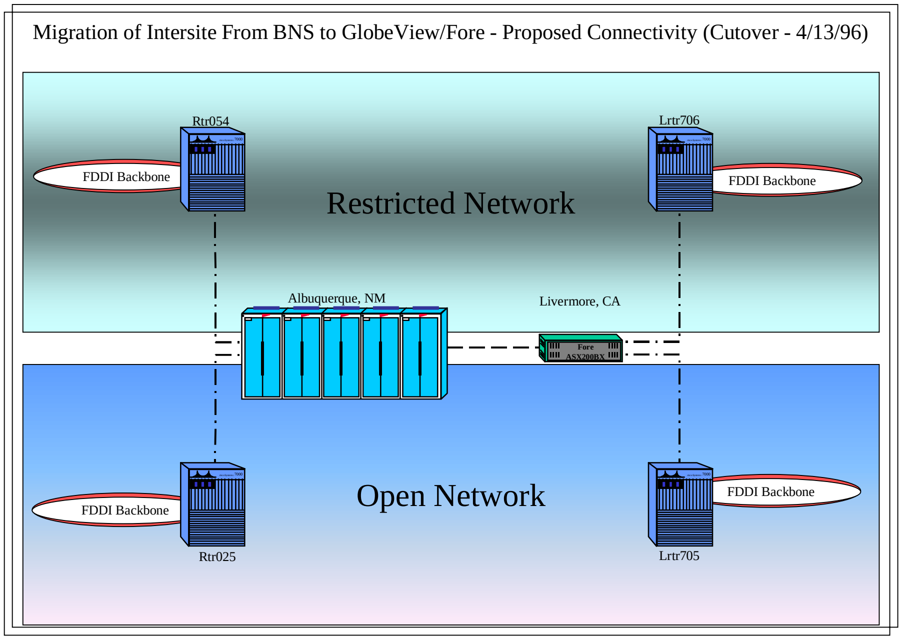

Figure 4-10. Diagram of *Completed* Phase-4 of the Migration Process.

Figure 4-10 shows the configuration of the Intersite Link after completion of Phase-4. The link utilizes the GlobeView via DS-3 over the private line to a FORE Systems ATM switch in California. We created sub-interfaces on routers in the EON and IRN in order to maintain routing between both campuses. We chose to keep this concept from the original Intersite link rather than bring the Livermore routers into our Albuquerque PVC mesh. If we are to include the Livermore routers into the Albuquerque PVC mesh, more configuration coordination between Albuquerque and Livermore networking personnel is required. In an effort to keep the configuration simple on both ends of the link, we decided to simply create a sub-interface on both sets of routers and create a single point-to-point PVC between them. See Appendix 4-1 for an example of the router AIP configurations.

The Migration Process: Phase-4 Lessons Learned

We added no new key techniques during Phase-4. We simply used our experience from previous phases in implementing these circuits.

We did have one lesson learned, though. In creating the PVC scheme for Rtr025 on the External Open Network (EON), we used a VCD / VCI which was already being used on a different VPI by external customers. The atm-pvc 175 6 175 aal5snap configuration command already existed on a different sub-interface on Rtr025. In implementing the Intersite link, we used 175 4 175. From a switch perspective, this appears valid.

However, from a router perspective, we used the same VCD. The router simply used the last VCD configuration, which broke the circuit for our external customers. Be sure to exercise caution when configuring PVCs on routers. Simply checking the PVC implementation from an ATM switch perspective was not enough in this case. Phase-4 of the migration process was completed April 13, 1996 and took 4 hours.

The Migration Process: Phase-5 External Open Network

The migration of the EON routers from the FDDI connectivity to the ATM switch was performed in a series of smaller phases or sub-phases. As stated earlier in this paper, we applied all lessons learned in the previous phases to this migration phase to improve the process. This section describes each of the sub-phases and presents any key techniques used or lessons learned during Phase-5 of the migration process.

Figure 4-11 illustrates the EON in its original configuration. Note this configuration is similar to the IRN original configuration.

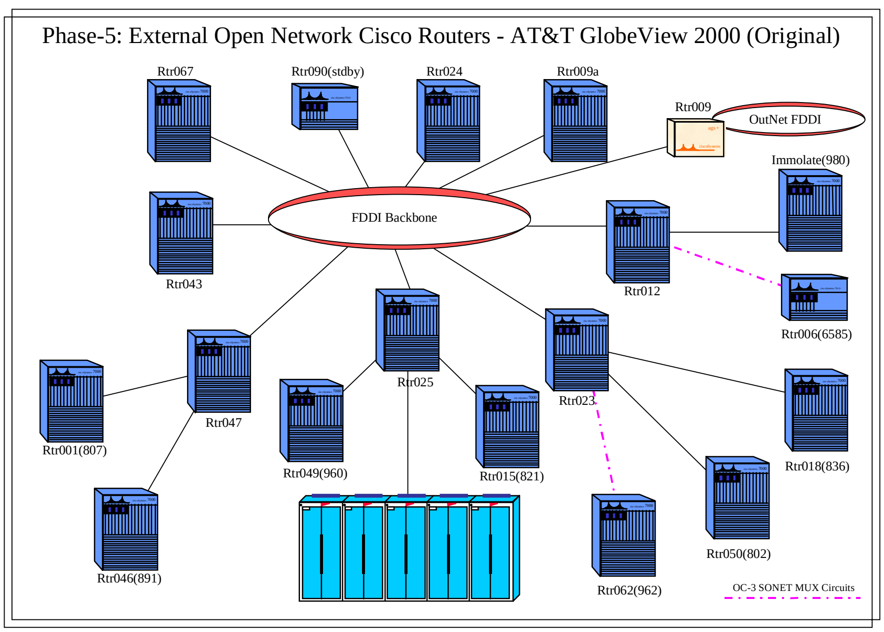

Figure 4-11. Diagram of Phase-5 of the Migration Process.

The EON network consisted of a central FDDI ring with distributed routers connected via SONET/DS-3 utilizing router HSSI interfaces and CSU/DSUs. Note the Rtr025 router is already connected to the GlobeView. This was done during Phase-2 of the IRN and provided routing for EON test customers.

The first cutover sub-phase included 3 routers, Rtr024, Rtr014, and Rtr025. The Rtr025 router (link between two technologies much like Rtr057 of the IRN) was already connected to the GlobeView, but did not exist on an EON production PVC mesh. We created the AIP sub-interfaces and configured the production mesh including each EON router migrated to the GlobeView. We had the circuits provisioned on the GlobeView in advance, thus making the cutover easier. We simply replaced the HSSI interfaces with AIP interfaces on the EON routers and configured the AIP interface as appropriate.

The second cutover sub-phase involved removing most functionality of the central FDDI ring. We had staged a new Rtr009 router (Rtr009a) to the GlobeView in preparation for this cutover. The Rtr009 router is between the EON production network and the OutNet (Internet). We use this router for key IP route filtering from the Internet Service Providers. We decided to move the function of Rtr025 as link between two technologies to Rtr009a. This prepared us for the next cutover, which included some relocating of routers and actually removing the EON Central FDDI ring.

During the third and fourth cutover sub-phases we completely removed the central FDDI ring. We no longer had the Cisco AGS+ Rtr009 router. The Cisco 7000 Rtr009 router was configured with the OutNet FDDI ring and a HSSI interface to a customer router in a different building (building 980). All that remained was the distributed routers. Refer to Appendix 4-2, 4-3, and 4-4 for illustrations of the Phase-5 sub-phases.

Our fifth and final cutover sub-phase proved extremely successful. With all techniques learned in the previous phases, we had no problems during this sub-phase. Figure 4-12 illustrates the final EON configuration.

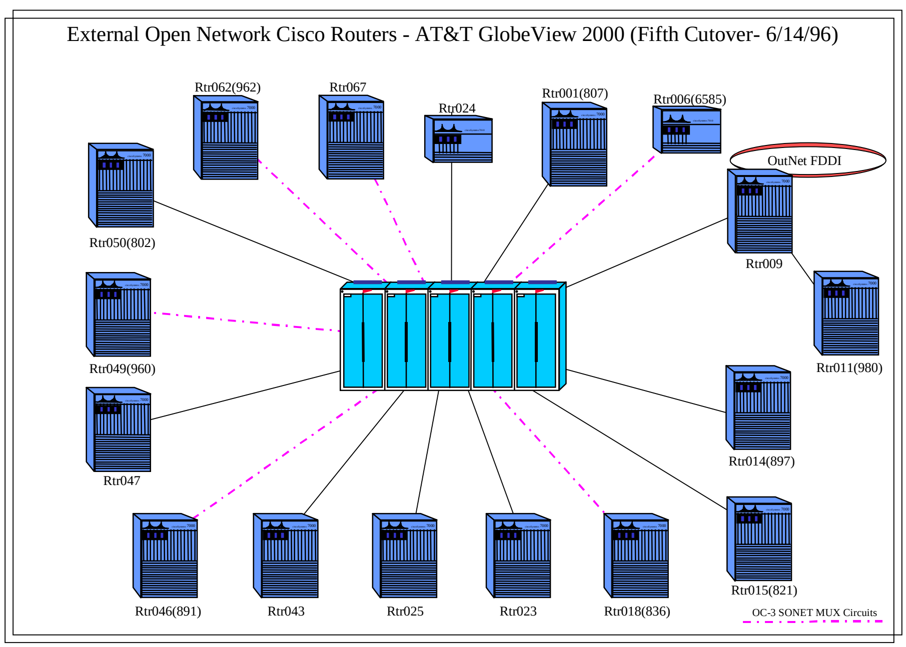

Figure 4-12. Diagram of *Completed* Phase-5 of the Migration Process.

Notice only Rtr011(980) remained on a HSSI point-to-point circuit. Since this is a remote customer's router, it was not yet considered a “production backbone router” and was migrated to the GlobeView at a later time.

The Migration Process: Phase-5 Key Techniques

To summarize Phase-5 techniques, all VP/VC circuits on the GlobeView were provisioned in advance. The physical circuits were built in advance, with proper attenuation, looped at the routers. We prepared configuration of the PVC mesh for the routers in advance. This included the entire mesh being applied to the router during the cutover sub-phase. As we progressed through the cutover sub-phases, this became a two-person operation. One person walked around the campus and simply powered off the router and replaced the HSSI interface with the AIP interface and connected the fiber for circuits back to the GlobeView. It became easy enough to identify which fiber strand was transmit vs. receive by just waving the fiber near the AIP ports and watching the lights on the AIP. (*Note: the appropriate cautionary measures need to be considered due to light emitted by the laser of the fiber optics.*) The other person had already entered the router's console via ISDN and immediately applied the configuration CUT/PASTE technique.

The Migration Process: Phase-5 Lessons Learned

We only had one lesson learned for Phase-5. We had been experiencing intermittent slow-down across our firewall router. Upon troubleshooting, we noticed that the Output Queue on the EON ATM interface would begin to fill and the router would eventually reset the interface. In troubleshooting on the GlobeView and the Cisco 7000 router, we determined the technicians were preparing the circuits for the next cutover sub-phase by pulling the loopback from the GlobeView port and extending it to a loopback at the next cutover sub-phase router. During this time, the GlobeView would pass Operation, Administration, and Maintenance (OAM) cells to each port associated with the VPI 2, the EON mesh. This caused the firewall router’s Output Queue to fill on the ATM interface for the EON. We have the ability of turning off this distribution of OAM cells, and have done so on the EON mesh-to-firewall circuits. We have performed several tests and have received a "maintenance" IOS release from Cisco in trying to resolve this problem.

During this last cutover sub-phase, we migrated four distributed routers in 30 minutes as a result of all the key techniques used and lessons learned from the previous phases.

Summary of Lessons Learned

A summary of the lessons learned during our migration of the Cisco routers to the GlobeView includes the following:

- OSPF Area cross-match
  If you use IP OSPF routing and choose a scheme such as ours, be careful with crossing areas. The router applies the configuration for OSPF on a line-by-line basis. We had inadvertently left the original OSPF Area 57 ahead of the OSPF Area 0 command in the configuration file.

- IP route-cache cbus on AIP sub-interface
  If you use fast-switching, be sure to apply the ip route-cache cbus command on the main AIP interface rather than the sub-interface. This will save you time.

- Customers making their local changes at same time
  Do not allow customers, or anyone else, make network changes during this significant migration effort. Control all network changes. *Just say - NO!*

- GlobeView upgrade and automation of VP/VC provisioning
  Be cautious of any upgrades to your core switches and if you choose automation, be aware of the consequences.

- Disabling router console logging
  I recommend you do not turn off the router console logging. Even though it may be difficult to accomplish your configuration changes, seeing those errors is STILL important.

- Router AIP microcode bug (AIP MC 10.8)
  Cisco IOS 10.2(8) Operating system with 10.8 AIP microcode has bugs. I recommend trying IOS 10.2(11) first, especially if you support all protocols such as we do.

- Cisco IOS 10.3(9) AppleTalk bugs
  If you are only running IP and considering IOS 10.3(9), keep in mind AppleTalk on the AIP has bugs. Test any IOS for as much functionality as possible.

- Consider ATM switch *AND* router interpretation of PVC configuration
  When you are creating the PVC schemes, be sure to evaluate from both a switch perspective as well as a router perspective.

- GlobeView passing OAM cells – configurable parameter
  Be aware of the OAM cells the ATM switches may pass if a port goes down. This may cause you some nightmares.

The Final Configuration

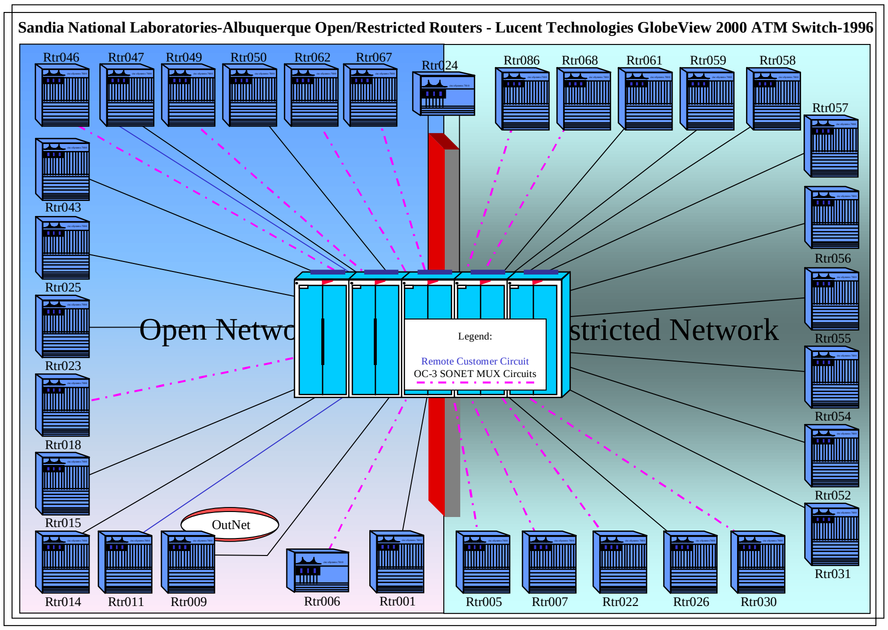

Figure 5-1. The Final Configuration – EON and IRN.

Figure 5-1 illustrates Sandia's External Open Network and Internal Restricted Network backbones with Cisco 7000 routers connected to the Lucent Technologies GlobeView 2000 ATM Switch. Note the different backgrounds represent the different network environments.

Conclusion

The project of migrating the Cisco 7000 series routers from FDDI interconnectivity to a core ATM switch began approximately February 1996. Before we moved the first pair of fiber optic cables, we had used several hours for configuration planning. We configured Cisco 7000 routers in a test bed and applied our knowledge and experience to simulate the production environment to the best of our resources. Testing the Cisco IOS version compatibility within any particular network environment is probably the biggest task, especially when your network supports most routable protocols such as Sandia’s Internal Restricted Network does. Much documentation was also generated in an effort to check/verify all configuration issues among all team members. As we began the migration process, we had our ambitions high and were eager to finish the process. We had initially planned migrating all routers to the core ATM switch in four large phases. As we progressed, we encountered some problems and changed our method of moving the routers to the core switch. We determined that moving the routers in several small cutover phases provided more time for configuring the devices in advance and verifying the configuration. When the maintenance time was scheduled, the actual cutover time was reduced, thus reducing network downtime for customers.

We have seen several valuable techniques utilized in the migration process. If we did not have ISDN connectivity to the router's console we would have certainly been doing much more walking than desired. With the abilities of X-Window CUT/PASTE, the configuration was simply applied to the ISDN router console window from files built before the individual migration phase. The method of addressing the devices based on the IP address of the device provides documentation within the configuration and makes troubleshooting easier. As the end of the migration process became reality in June 1996, we had collected some valuable information from the lessons we learned. It was the experience gained by those lessons that we were able to move the last four routers to the core ATM switch in 30 minutes.

Acknowledgements

In order to accomplish a successful project such as the migration of production network backbone routers from FDDI interconnectivity to the GlobeView ATM switch, several highly talented personnel were involved. We had teams from four different organizations within Sandia working hard together. The list of personnel includes:

SNL, Albuquerque:

4616 Advanced Networking Integration	Steven Gossage, John Naegle, Tan Chang Hu, Luis Martinez, Joseph Brenkosh, John Eldridge
4417 Telecommunications Operations-II	Bruce Whittet, David Evans, Spencer Nelson, Frank Quintana, Pat Torrez, Vicki Williams
4914 Telecommunications Operations-I	George Rivera, George Yonek, Frank Castelluccio, Debi Brunty, Joseph Maestas, All Tech. Control Center Personnel

SNL, Livermore:

8930 Communications and Network Systems	Rich Gay, Frank Bielecki, Joe Brazil

I would like to thank the following individuals for their team-oriented contributions to the construction of this document:

Joseph Brenkosh, David Evans, Tan Chang Hu, Luis Martinez, Spencer Nelson, Frank Quintana, Lawrence Tolendino, Pat Torrez, Bruce Whittet, and Vicki Williams

References

[1] The ATM Forum Web Site http://atmforum.com. The ATM Forum is an international non-profit organization formed with the objective of accelerating the use of ATM (Asynchronous Transfer Mode) products and services through a rapid convergence of interoperability specifications. In addition, the Forum promotes industry cooperation and awareness. The ATM Forum consists of a worldwide Technical Committee, three Marketing Committees for North America, Europe and Asia-Pacific as well as the User Committee, through which ATM end-users participate. Worldwide Headquarters, 2570 West El Camino Real, Suite 304 - Mountain View, CA 94040-1313, +1.650.949.6700 Phone , +1.650.949.6705 Fax

[2,3] Cisco Connection Online Web Site http://www.cisco.com. Cisco Systems, Inc. (NASDAQ: CSCO) is the worldwide leader in networking for the Internet. Cisco products include routers, LAN and ATM switches, dial-up access servers and network management software. These products, integrated by the Cisco IOS software, link geographically dispersed LANs, WANs and IBM networks. Cisco Systems news and product/service information are available on the World Wide Web site at http://www.cisco.com. Cisco Systems is headquartered in San Jose, Calif.

Glossary of Acronyms

AIP – ATM Interface Processor
AIP MC – ATM Interface Processor MicroCode
AT – AppleTalk
ATM – Asynchronous Transfer Mode
CA - California
CLNS – ConnectionLess Network Service
CON – California Open Network
CSU – Channel Service Unit
CW – CiscoWorks
CWSI – CiscoWorks for Switched Internetworks
DECNET – Digital Equipment Corporation NETwork
DS-3 – Digital Signal (level 3)
DSU – Data Service Unit
EON – External Open Network
FDDI – Fiber Distributed Data Interface
GDC – General Datacom Corporation
HP – Hewlett-Packard
HSSI – High Speed Serial Interface
IOS – Internet Operating System
IPX – Internet Packet eXchange
IRN – Internal Restricted Network
ISDN – Integrated Services Digital Network
ISS – Interface and Services Stage
LAN – Local Area Network
LANE – LAN Emulation
LS1010 – LightStream 1010
LTP – Logical Termination Point
MAN – Metropolitan Area Network
NM – New Mexico
NMS – Network Management System
OAM – Operation, Administration and Maintenance
OC-3 – Optical Carrier (level 3)
OSPF – Open Shortest Path First
PVC – Permanent Virtual Circuit
RAN – Restricted Access Network
RME – Resource Manager Essentials
SMDS – Switched Multimegabit Data Service
SMF – Shared Memory Fabric
SNCP – Service Node Control Processor
SNL – Sandia National Laboratories
SONET – Synchronous Optical NETwork
TCP/IP – Transport Control Protocol/Internet Protocol
VC – Virtual Circuit
VCD – Virtual Circuit Descriptor
VCI – Virtual Channel Identifier
VP – Virtual Path
VPI – Virtual Path Identifier
WAN – Wide Area Network
XNS – Xerox Network Service

Appendix

A-1.1: The Internal Restricted Network Original Configuration.

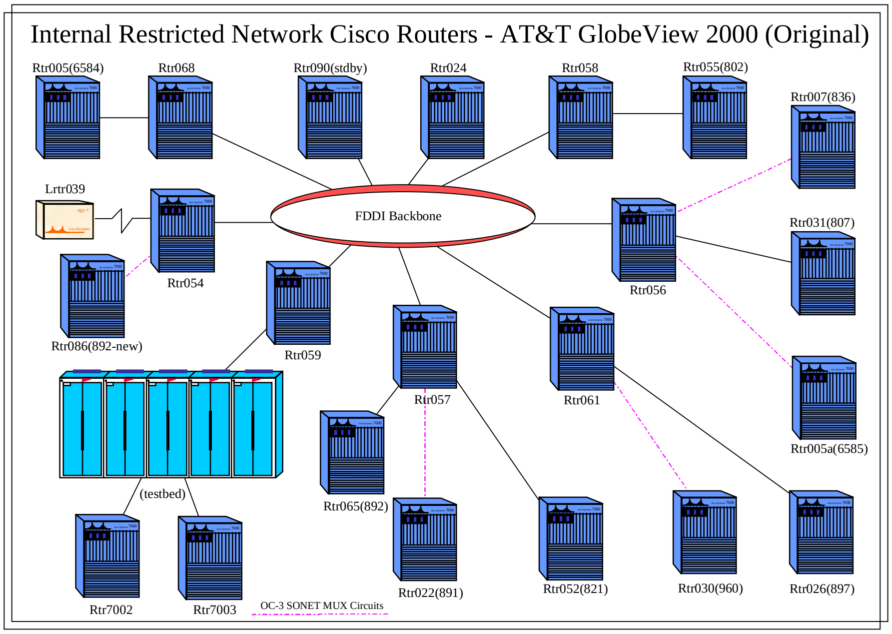

A-2.1: CiscoView Graphical User Interface Display

A-2.2: CiscoWorks 2000 VLAN Director Application Display

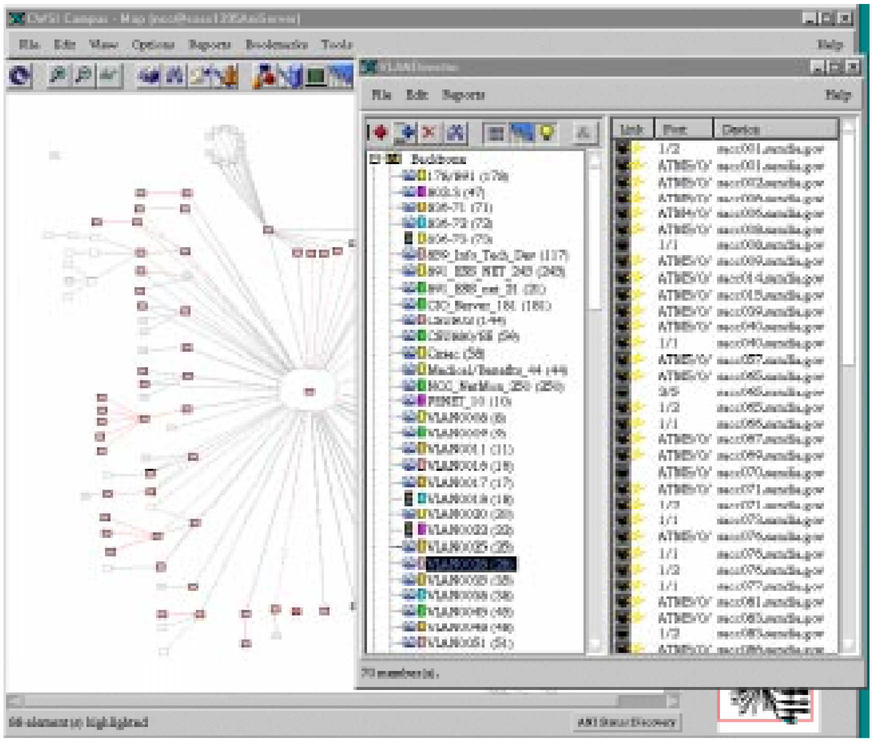

A-2.3: CiscoWorks 2000 ATM Director Application Display.

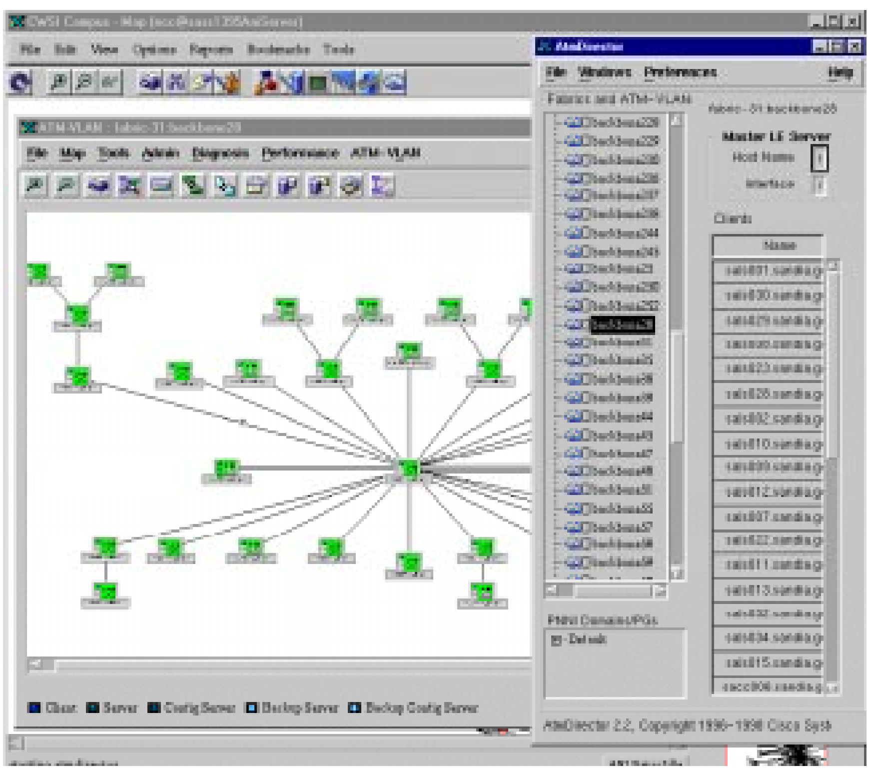

A-2.4: CiscoWorks 2000 Traffic Director Application Display

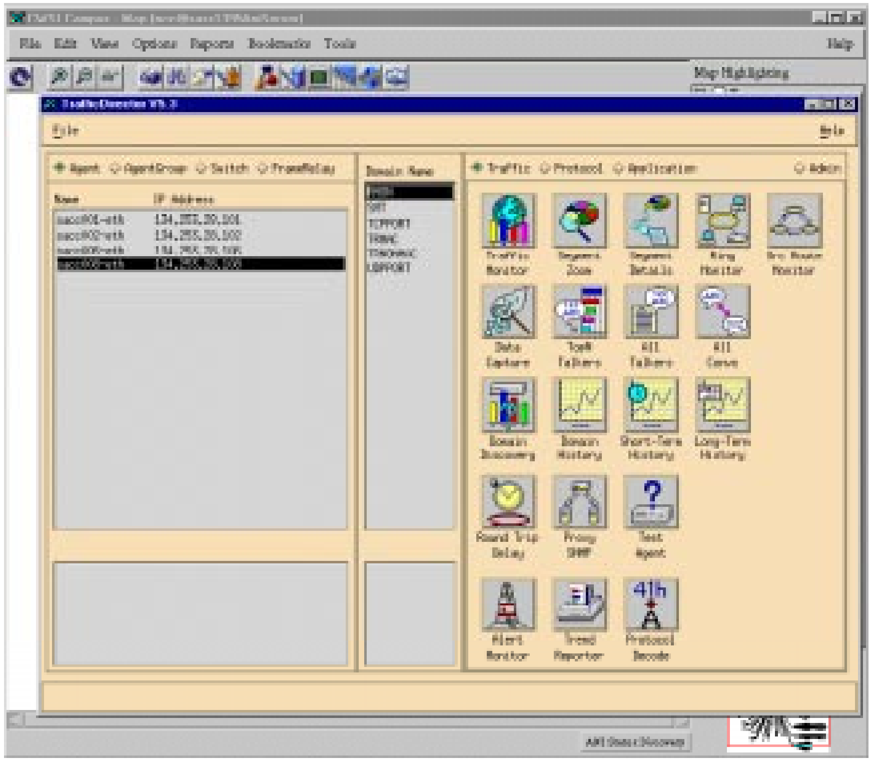

A-2.5: CiscoWorks 2000 Resource Manager Essentials Display

A-3.1: Mapping of the Phase-1 Routers – BEFORE the Migration

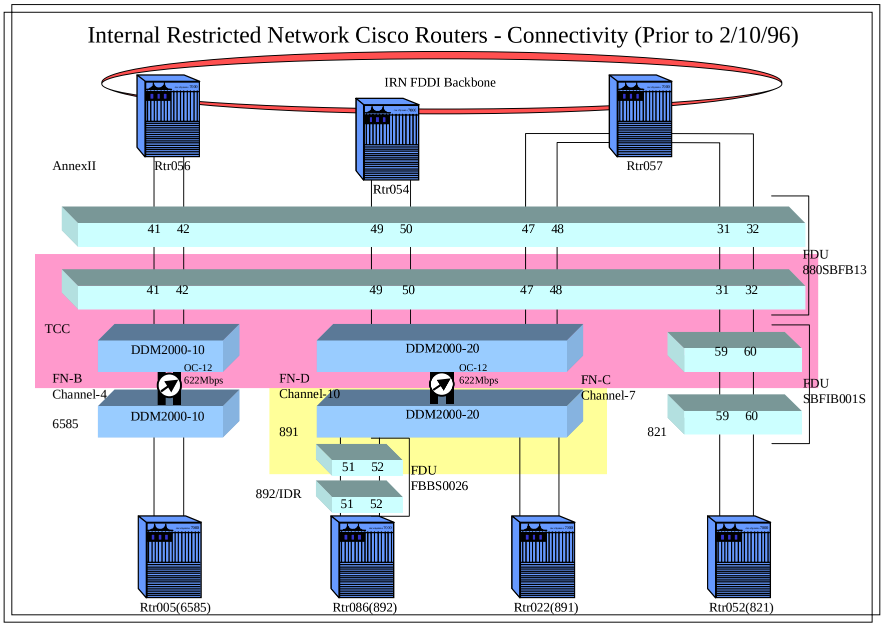

A-3.2: Mapping of the Phase-1 Routers – AFTER the Migration

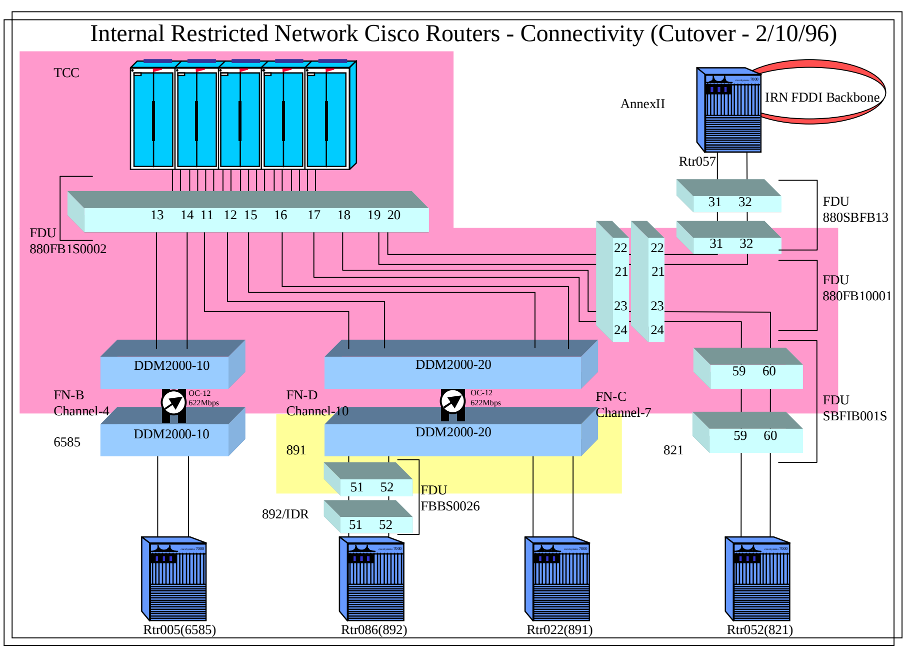

A-4.1: Sample ATM Configuration for the Intersite Link

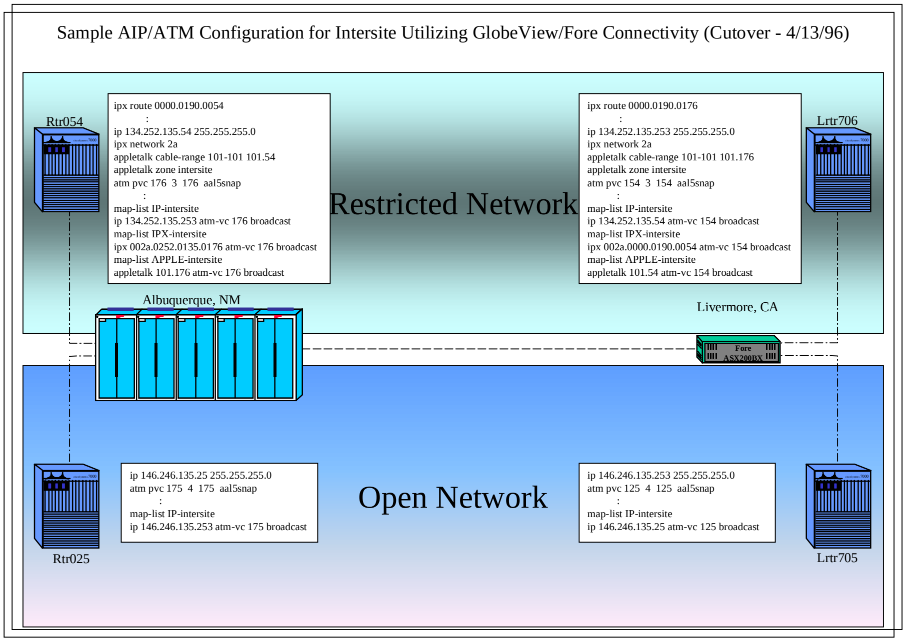

A-4.2: Mapping of the Phase-5, First Cutover Sub-phase

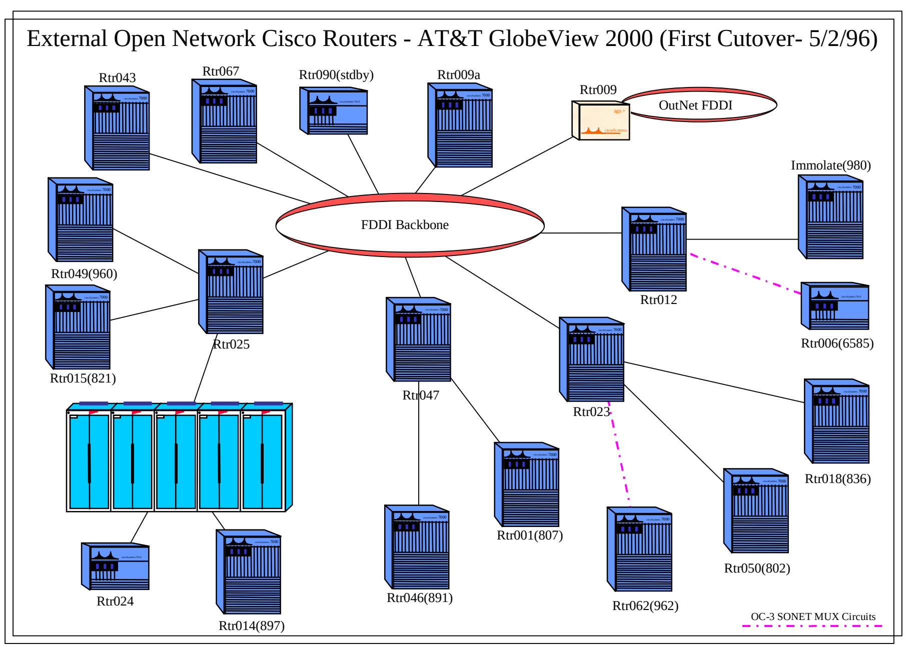

A-4.3: Mapping of the Phase-5, Second Cutover Sub-phase

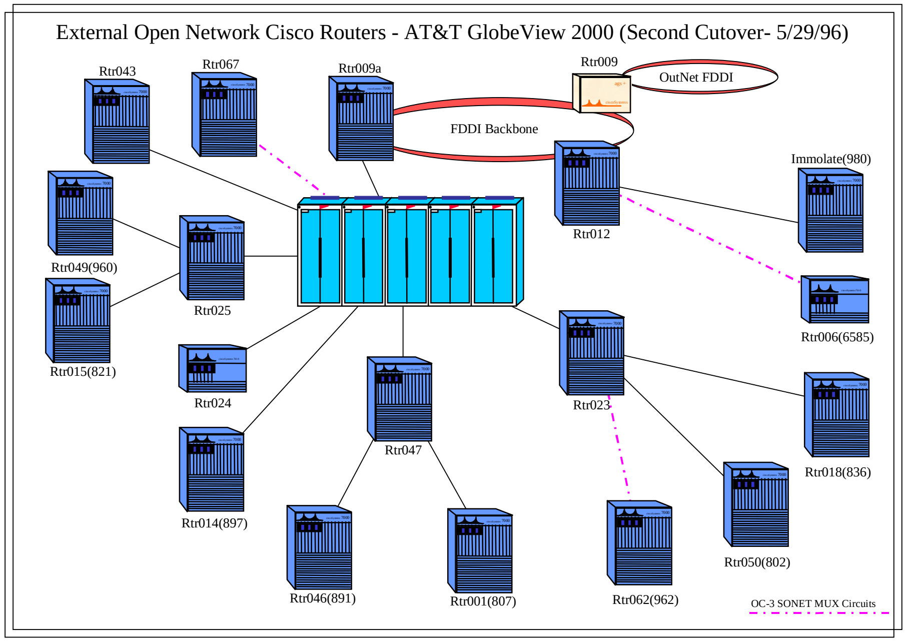

A-4.4: Mapping of the Phase-5, Third and Fourth Cutover Sub-phases

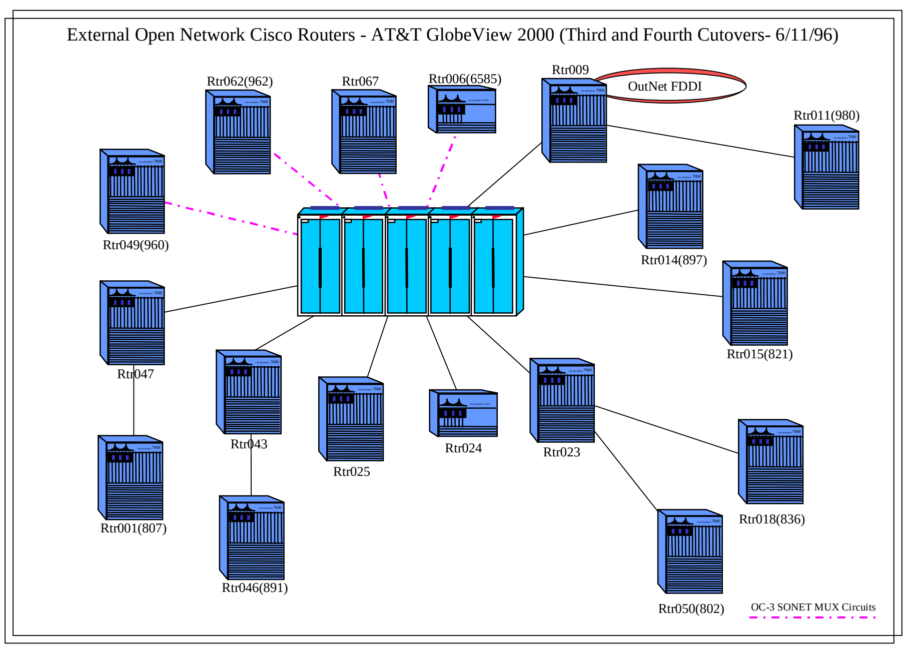

DISTRIBUTION:

1	MS 0469	J.F. Jones, Jr., 4600
1	MS 0801	M.J. Murphy, 4400
1	MS 0803	L.H. Pitts, 4900
1	MS 0806	D.C. Jones, 4401
1	MS 0806	M.O. Vahle, 4616
1	MS 0806	J.P. Brenkosh, 4616
1	MS 0806	J.M. Eldridge, 4616
1	MS 0806	S.A. Gossage, 4616
1	MS 0806	T.C. Hu, 4616
1	MS 0806	L.G. Martinez, 4616
1	MS 0806	J.H. Naegle, 4616
1	MS 0807	I.C. Alexander, 4417
1	MS 0807	D.P. Evans, 4417
1	MS 0807	S.D. Nelson, 4417
1	MS 0812	M.R. Sjulin, 4914
1	MS 0812	R.L. Adams, 4914
7	MS 0812	T.V. Archuleta, 4914
1	MS 0812	F.J. Castelluccio, 4914
1	MS 0812	J.A. Chavez, 4914
1	MS 0812	S.B. Coningham, 4914
1	MS 0812	E.J. Klaus, 4914
1	MS 0812	J.H. Maestas, 4914
1	MS 0812	P.L. Manke, 4914
1	MS 0812	G. Rivera, Jr., 4914
1	MS 0812	L.F. Tolendino, 4914
1	MS 0812	B.C. Whittet, 4914
1	MS 0812	V.K. Williams, 4914
1	MS 0812	G.A. Yonek, 4914
1	MS 1012	P.M. Torrez, 4811
1	MS 9003	D.L. Crawford, 5200
1	MS 9012	S.C. Gray, 8930
1	MS 9012	F.T. Bielecki, 8930
1	MS 9012	J.G. Brazil, 8930
1	MS 9012	R.D. Gay, 8930
1	MS 9018	Central Technical Files, 8940-2
2	MS 0899	Technical Library, 4916
1	MS 0619	Review & Approval Desk, 15102 For DOE/OSTI